PATENT APPLICATION

for

Optically Switched Communication Network

The present invention relates to communication networks and in particular to large optically switched communication networks.

BACKGROUND OF THE INVENTION

High date rate optical communication systems are well known. The communication industries have developed standards to facilitate interconnectivity among various communication suppliers. One popular standard utilizing time division multiplexing (TDM) to transmit information at high data rates through optical fibers is OC-48/STM-16 which provides data rates at 2.5 Gbps. Other standards at higher data rates are available but OC-192/STM-64 at 10 Gbps represents a current practical upper limit for TDM, with same result being done at 40 Gbps. Higher data rate can be provided with wavelength multiplexing in which data is multiplexed according to wavelength and time. Modern systems can routinely pack 40 10-Gbps channels through a single fiber for an aggregate bit rate of 400 Gbps. This is referred to as dense wavelength division multiplexing (DWDM).

The International Telecommunication Union has established a standardized separation grid calibrated on a krypton spectral line at 193.1 THz (wavelength equal to 1.57 micron) ranging from 186.0 THz to 201.0 THz which provides a total available bandwidth of 15,000 GHz. This bandwidth range is separated into channels which are separated by 100 GHz, although some players squeeze in more channels by providing a separation of only 50 GHz. Separating the 15,000 GHz at 100 GHz spacings would provide 150 channels and at 50 GHz spacings about 300 channels could be provided.

The typical DWDM system includes (1) multiplexers for both time and wavelength for combining signals for transport, (2) demultiplexers for disassembling the aggregate signal so

that each signal can be delivered to the appropriate receiver, (3) active or passive switches and/or routers to direct each signal component, (4) filters to provide wavelengths selection and (5) optical add-drop multiplexers which give the service provider the ability to pick up and drop off individual wavelength components at intermediate locations throughout the network. DWDM is important because expanding a network by putting more fiber in the ground costs about $70K per mile. Adding bandwidth using DWDM costs about one-sixth this amount.

The principal form of optical switching is nothing more than a sophisticated digital cross-connect system. In the early days of data networking, dedicated facilities were created by manually patching the end points of a circuit at a patch panel, thus creating a complete four-wire circuit. Beginning in the 1980's, digital cross-connect devices such as AT&T's Digital Access and Cross Connect (DACS) became common, replacing the time-consuming, expensive, and error-prone manual process. Most current switches convert signals from incoming optical fiber to electronic signals, switch electronically and then convert the signal back to optical for transmission into the output optical fiber. These switches are referred to as O-E-O switches. The digital cross-connect is a simple switch, designed to established long-term temporary circuits quickly, accurately and inexpensively.

Enter the world of optical networking. Traditional cross-connect systems worked fine in the optical domain, provided no problems were happening in the O-E-O conversion process. This, however, was one of the aspects of optical networking that network designers wanted to eradicate from their functional requirements. Thus was born the optical cross-connect switch. The first of these to arrive on the scene was an optical switch provided by Lucent Technologies. The switch was based on a switching technology called the Micro-Electrical Mechanical System (MEMS); the switch was the world's first all-optical cross-connect device.

MEMS relies on micro-mirrors. The mirrors can be configured at various angles to ensure that an incoming optical beam strikes one mirror, reflects off a fixed mirrored surface, strikes another movable mirror, and is then reflected out an egress fiber. These devices are now

commercially deployed and offer speed, a relatively small footprint, bit rate and protocol transparency, nonblocking architecture, and highly developed database management. Fundamentally, these devices are very high-speed, high-capacity switches or cross-connect devices. They are not routers, because they do not perform layer-three functions.

One of the main difficulties associated with previously proposed all-optical networks is the need for DWDM colors to be changed as they pass through the core switching network. The need comes from traditional switching architectures such as Clos-type to "scramble" or rearrange subchannels in order to achieve full connectivity.

What is needed is a better communication network providing on a national scale optical switching without wavelength conversion.

SUMMARY OF THE INVENTION

The present invention provides a large communication network suitable for nationwide or worldwide utilization. A plurality of area code nodes are connected with all-fiber-optic links with all-optical switches. A routing algorithm provides one or more communication links from each area code node to every other area code node so that information never has to change carrier wavelength as it travels the network. Each area code node contains circuits that are provided to connect individual users to the network.

A first preferred embodiment of the present invention is a nation-scale high-speed communications network. The network provides dedicated, circuit switched connections between end users, and can support 100 million user nodes with 25 MHz of full duplex bandwidth per user. This amount of bandwidth can support 100 Mb/s data capability and/or multiple television channels. The method aggregates 15,000 MHz of bandwidth onto each of 300 different wavelengths at 50 GHz spacings using electro optical modulation and sub-frequencies spaced at 4 GHz. The standard link provides 1200 different fiber colors (consisting of 300 wavelengths on 4 fibers). Links are routed between 250 area code nodes

using all-optical switches. Using a special routing algorithm applicants call "magic square," data never has to change carrier wavelengths as it traverses the network.

DETAILED DESCRIPTION OF PREFERRED EMBODIMENT

First Preferred Embodiment

A top level description of a first preferred embodiment is shown in FIG. 1. In the discussion which follows, it should be understood that all quantities refer to a specific example of how the network might be implemented in a self-consistent manner, but that any specific numbers could be modified as the result of a more detailed analysis for a specific real world implementation. Here, for the first preferred embodiment, we consider a nation-scale network with end users assigned to one of 250 area codes with roughly equal numbers of end users in each. For instance, area code #1 has been assigned to San Diego, #40 to Seattle, #200 to Washington D.C., and #240 to some subset of international users. The proposed network can have about 400,000 User Nodes per Area Code. As seen in FIG. 1, Optical Cross-Connect switches (OXC's) 2 associated with each area code are located at mesh nodes 4 tied together in a mesh network which allows switching of optical signals from any particular area code to any other area code. The particular mesh network would make maximum use of intercity fiber trunk lines which have already been installed.

In this preferred embodiment all long haul communication is through optical fibers operating in the wavelength range centered at about 1.57 micron (corresponding to about 193.1 THz). The network is designed to operate at frequencies between 186 THz to 201 THz for a total bandwidth of 15,000 GHz. At 50 GHz spacings, this provides 300 "color" channels per optical fiber. Four separate fibers provide a total of 1200 communication channels.

There are in this embodiment 1200 separate wideband communications channels (each with 15 GHz of useable optical bandwidth) leaving and entering each area code. We call these wideband channels FiberColors, and they are distributed as 300 different DWDM wavelengths (standard 50 GHz spacing) on 4 separate fibers. The preferred optical network

operates in the C and L bands at a center frequency at 1570 nm (193.1 THz). Thus there are 8 lit fibers between an Area Code and its corresponding OXC, 4 for outgoing traffic and 4 for incoming traffic as shown at 6 on FIG.1.

If we divide 1200 FiberColors by the number of area codes (300) we have an average of 4.8 FiberColors per area code. However, the 1200 outgoing FiberColors from any particular area code (say San Diego) are allocated based on usage demand to the 250 area codes with a total bandwidth of ~15 THz per fiber. For instance, the FiberColors for traffic from San Diego might be allocated at a particular time as follows: 10 for traffic to Washington, 6 for traffic to Seattle, 1 for traffic to Atlanta, etc., until all 1200 FiberColors are accounted for. It is expected that the actual allocation will be automatically adjusted periodically as demand shifts with time of day and day of the week. Therefore, at any particular time, the OXC switches must be configured so that each FiberColor from each origination area code is guided through the network to its destination area code without interference. (That is, the same fiber cannot be used simultaneously for two FiberColors operating at the same wavelength). It was not immediately obvious that this could be done, but applicants have developed an algorithm for accomplishing this task which appears to be robust and to converge in a practically short time. We call this algorithm a Magic Square Algorithm, because the underlying matrices of FiberColors which need to be allocated have rows and columns which add up to the same number. This problem of allocating FiberColors along with its solution is discussed in detail in the section of this specification entitled "Magic Square Software." The solution of this problem is a key technical innovation, as it enables the deployment of a nation-scale all-optical network with a relatively small number of channels without the disadvantages of having to convert any optical signal to an electrical signal or to another DWDM wavelength between the data source and the data destination area codes.

Operation of the OXC Switches 2 at the Mesh Nodes 4 is discussed in more detail in the section of this specification entitled "Optical Cross Connect Switches." In preferred embodiments, all DWDM wavelengths are demultiplexed before optical switching, and then remultiplexed after switching. No wavelength separation is required at a resolution finer than

the standard 50 GHz DWDM spacing, so that standard components can be used. (Finer channel resolution only occurs within the source and destination area codes). Customized switches which combine wavelength separation with the optical switching may also be possible. Optical Amplifiers (such as erbium doped fiber amplifiers, EDFA's) are used throughout the network as necessary to maintain appropriate optical signal strength.

Area Code Operation

The implementation of a preferred area code network 8 at the Area Code level can be understood with reference to FIG. 2. Two main functions happen at the area code level. The first main function, which can be geographically distributed, is the passive splitting and combining of the four outgoing and incoming fibers to 100 different community nodes CN 1 to CN 100 as shown on FIG. 2. That is, each of the four incoming fibers to the area code is passively split into 100 different fibers going to different community nodes within the area code. Fibers from the 100 different community nodes are passively combined into the outgoing fibers. The minimum 20 dB loss associated with this process is compensated with optical amplification in the appropriate places. With this passive splitting and combining, the full 17.28 THz of bandwidth available in the area code is accessible at each community node. (Of this, up to 10 THz might be in use at one time, with 400,000 users having access to 25 MHz each).

The second main function implemented at the Area Code level (preferably with master signals generated at the national level) is the creation of wavelength reference combs defining optical subfrequencies at a finer resolution than the standard DWDM grid. This reference comb, shown in FIG. 5 and discussed in more detail later, consists of narrowband optical reference signals with a 4 GHz spacing, stable on an absolute scale to about 500 MHz. Each of the 300 DWDM wavelengths separated by 50 GHz has 6 optical subfrequencies associated with it as shown in FIG. 5. There are different ways to generate the reference frequency comb which are well known. Multiple stabilized lasers could be used and combined passively (with optical amplification), or fewer multi-frequency fiber lasers with appropriate etalons could be used. The frequency comb is generated in the area

code and distributed from there to the 100 community nodes to form the basis of carrier optical frequencies onto which the data is modulated.

Community Node Operation

The community nodes contain the critical racks of equipment which modulate user data onto the optical subfrequency carriers in the frequency comb, and demodulate data from the optical carriers for distribution to the end users. As shown in FIG. 3, each community node (CN) has access to all of the 1200 upstream and downstream FiberColors (through passive splitters and combiners), and connects to 25 Neighborhood Nodes NN 1-25. The modulation and demodulation method is discussed later with reference to FIGS. 6 and 7. The community node takes upstream analog signals originating from different users and modulates them onto 25 MHz of spectrum within a particular FiberColor depending on the destination area code of the data. Similarly, the community node takes downstream optical signals, and demodulates 25 MHz of spectrum for forwarding on to a designated end user. Because of the distances involved, the user data is generally transmitted to and received from the Neighborhood Nodes over optical fibers, where Coarse Wave Division Multiplexing (CWDM) can be used as a cost effective solution. However, the user bandwidth is in the form of an analog oscillating electric voltage at the interface with the Community Node rack of modulators and demodulators. Up to 4 GHz of bandwidth is transmitted to or received from each Neighborhood node at a given time, allowing for 160 users at 25 MHz each. Extra CWDM fibers are used to carry broadcast video information to the Neighborhood Nodes for shared use among multiple users. The details of the community node are displayed in FIG. 3.

Neighborhood Node Distribution

The neighborhood node (shown in FIG. 4) is basically an analog collection hub. End users (such as houses or employees in an office) are connected to analog coaxial cables in groups of 16 as shown at 10 in FIG. 4. The cables are operated using 900 MHz of bandwidth, with 25 MHz per user downstream, 25 MHz per user upstream, and 100 MHz of shared broadcast video bandwidth (with appropriate isolation bands). This allocation would be modified as

necessary to make maximum use of already installed cable systems. The signals from 10 cables in each Neighborhood Node (serving 160 end users) would be combined onto fiber optic cables for communication with the Community Node using standard CWDM techniques. Each end user could basically use their 25 MHz of analog bandwidth as they pleased, paying for end-to-end circuit switched connections to the user with which they were communicating. Connection to multiple simultaneous data sources or recipients is handled by division of the user bandwidth at the Community Node before modulation onto the optical carrier.

Optical Subfrequency Comb

A critical part of the preferred ColorFast network is the generation of the accurate frequency reference comb 20, shown in FIG. 5. This allows the precise placement and retrieval of 25 MHz user bandwidth channels. After generation of the subfrequency reference, relative frequency offsets are used, so that the components themselves do not have to have accurate absolute wavelength calibrations, which would be expensive. The reference subfrequency comb is either built up from separate modulated, filtered and amplified DWDM lasers, or from fiber lasers. By doing this at the Area Code level, the expense is distributed over many end users. The comb 20 includes 6 subfrequencies (separated by 4 GHz) for each DWDM wavelength range (separated by 50 GHz) as shown in FIG. 5. Modulation and demodulation of data at the DWDM wavelength level is only done at the Community Nodes, so that the Network Optical Switches at these wavelength ranges are not required. User data is modulated directly into the spaces between the reference subfrequencies using electro-optic modulators. Some part of the reference comb is transmitted across the network along with the user data, helping to provide frequency calibration at the receiver electronics. Each of the 6 optical subfrequencies is associated with 2.4 GHz of bandwidth for user data. The bandwidth is limited to allow the use of lower cost silicon electronics, rather than requiring more expensive, higher bandwidth materials.

User Data Modulation
(From User to Network)

The process of modulating data from the end user onto a particular optical subfrequency is shown in FIG. 7. The optical signal originating from a reference fiber (after appropriate optical amplification) is passed through a narrowband tunable filter 40 to isolate the desired subfrequency. (The tunable filter can be calibrated by sweeping over the subfrequency comb). The isolated narrowband frequency source is sent through an arrangement of electro-optic modulators 42. The modulation source is obtained from the end user by passing the upstream RF cable signal through an RF filter to isolate the 25 MHz of user data and passing that through a frequency converter so that it is at the frequency offset (designated by the network control for the particular end to end data transmission session) as shown at 44. The output of the EO modulator arrangement contains some residual reference signal, along with the 25 MHz of user data at the appropriate optical frequency offset. Details of how the electro-optic modulator achieves this result while minimizing extraneous sidebands are given below in a section entitled ____________. This is a critical technology for the proposed ColorFast network. The different modulator outputs are then combined passively (with appropriate filtering and optical amplification) resulting in a broadband signal to be transmitted across the network on a particular FiberColor.

User Data Demodulation
(From Network to User)

FIG. 6 shows details of how the preferred ColorFast network transfers a FiberColor from the origination area code to the destination area code. Embedded in that FiberColor are six 2.4 GHz data channels 28 for multiple recipients in the destination area code. Demodulators at the Community Nodes can be tuned to retrieve the appropriate 2.4 GHz channel 28A. Demodulators can be shared among different end users in the same Neighborhood Node. First, data from one of the four fibers is sent along with the signal from a reference fiber through a narrowband tunable filter 30. (The four fibers can either be switched to go through the appropriate demodulator, or separate demodulators can be designated for different fibers.)

The tunable filter contains both broadband and narrowband etalons to select a particular desired 4 GHz subfrequency band. After attenuating the reference signal to match the intensity of the received data, the reference and data optical signals are combined on a broadband detector where they interfere. (In practice a pair of detectors might be used to analyze an in-phase and out-of-phase component to get better noise performance). The electrical output of the detector circuitry 32 contains the 2.4 GHz of user data associated with the particular subfrequency. This is sent through a tunable RF filter 34 to separate out the desired 25 MHz of end user data. The carrier frequency of the data is then shifted to the allocated cable frequency of the end user, and combined with other end user data from other modulators before being sent to the Neighborhood Node and on to the end users as shown at 36.

Mesh Node

In the preferred ColorFast network architecture mesh nodes are responsible for providing all traffic routing and switching. The mesh nodes receive configuration commands from the computers that calculate the routing paths. Mesh nodes are located in the interior of the network as shown at 4A in FIG. 1 and at the edge of the network as shown at 4 in FIG. 1. The nodes act upon these commands, reprovisioning the network to provide bandwidth where it is required. The mesh nodes accomplish this provisioning and routing using all-optical cross-connects that are independent of protocol and encoding schemes. The routing computers are not necessarily contained in the mesh node although they may be. If not inside the mesh node, the mesh node receives routing commands from elsewhere in the network.

Mesh nodes in the interior of the network are responsible for switching colors to the proper location, and for providing some local add-drop and channel monitoring capability. Mesh nodes at the edges of the network provide an interface between the network and local area code traffic.

An example of a simplified mesh node is shown in FIG. 8. It has two mesh-node to mesh-node connections 52 and one area code or local connection 50. For simplicity each

connection contains only two input and output fibers instead of the 4 input and output pairs which is the baseline for the preferred architecture, and each fiber contains only two colors instead of the 300 envisioned for the actual network. Immediately after entering the node each input fiber is passed through a passive optical demultiplexer that breaks the incoming light into the 300 50-GHz color bands. Each color from all the input fiber is then input into an optical cross-connect switch and routed to a passive optical multiplexer. The optical multiplexer combines the 50 GHz color bands onto the output fiber which carries them to the appropriate area code or next network node.

The number of connections that can be supported by the mesh nodes is only limited by size of the optical cross-connect switch. Since each fiber is separated into it constituent colors before switching, the number of input and output ports which require switching grows quite rapidly. For example, a node having 3 connections each with 4 fibers and 300 colors per fiber requires switching for 3600 input and output ports. However, notice that since each input or output fiber can only carry one channel of each color, the switch does not need to provide full connectivity between all input and output ports. Thus, for the example given above, 300 small 12x12 cross-connects provide the required switching capability. For a node having 8 connections with four fibers per connection (32 fibers) and 300 colors per fiber each input color channel can be routed to one of the 32 output fibers. Thus, a large switch can be used which only provides connectivity between an input port and its 32 nearest neighbors. Alternatively 300 smaller 32x32 port switches could be used. Due to this reduced connectivity requirement a extremely large mesh node with 32 connections could be supported by a switch of moderate size such as a 256x256 MEMS based optical cross-connect.

The mesh nodes require a large number of optical multiplexing and demultiplexing units. However, this preferred ColorFast architecture does not require the development of devices different from those currently being made for telecommunications applications. The ColorFast requirement for demultiplexers and multiplexer are functionally equivalent to currently available 50 GHz DWDM filters. Several currently available technologies can

provide suitable channel separation including thin-film technology, diffraction gratings and arrayed waveguides.

Mesh nodes will also be convenient places to accomplish network grooming and signal amplification. For the sake of clarity this functionality has not been included in the initial description of the mesh nodes because it is not strictly required for switching. But in the deployed nodes signal grooming and amplification will preferably be included in each node in varying degrees depending upon the transmission requirement between the node's connections.

Calculating the required resources for a mesh node of any size is relatively easy. Refer to FIG. 9 which shows a generalized mesh node. The total number of connections to a particular mesh node is expressed as $N = N_M + N_L$, where N_M is the number of mesh node connections and N_L is the number of local connections. So if there are F_i fibers at the i^{th} mesh node to mesh node connection and L_j fibers at the j^{th} local node connection, the total number of input and output fibers, F_T at the node are

$$F_T = \sum_i^{N_M} F_i + \sum_j^{N_L} L_j .$$

Thus F_T multiplexers are required, F_T demultiplexers are required, and if n_c is the number of colors per fiber then the total number of ports X_T required by the optical cross-connect is

$$X_T = n_c F_T .$$

Neighborhood Node Modulator

FIG. 10 is a drawing of a neighborhood node modulator shown as 42 on FIG. 7. This preferred modulator is comprised of two Mach-Zehnder interferometers placed in series. Mach-Zehnder interferometers are fairly common in fiber optics communications systems and are often used as external modulators for transmission of data. They are often implemented as waveguides in a Lithium Niobate substrate and commonly have one leg with optical path length that can be varied by applying a voltage across the Lithium Niobate crystal. They are suitable for use with high frequency communications systems and can be

modulated at bandwidths in excess of 10 Ghz. Since the modulation proposed for the ColorFast system is more sophisticated than that required in current communications systems two interferometers are required to accomplish the required frequency shift and modulation. It is also easiest to implement the ColorFast scheme if both legs of the interferometers are fabricated such that they are independently variable.

FIG. 10 shows a schematic of the proposed modulator. The reference color is input into one leg of the first Mach-Zehnder interferometer. A mode coupler is used to distribute half of the input light to each of the legs of the interferometer. The optical path lengths of the two legs of the interferometer are controlled by two input signals. The two output ports of the interferometer are then passed through a second complimentary interferometer whose legs are varied by two different input signals. If done correctly the output of this device is a communications signal with data modulated onto a color-shifted carrier.

The voltages for each leg of the modulator are given below,

$$\begin{aligned} V_1 &= a(t)Cos(\omega_m t) + V_o \\ V_2 &= a(t)Sin(\omega_m t) + V_o \\ V_3 &= a(t)Sin(\omega_m t) + V_o \\ V_4 &= -a(t)Cos(\omega_m t) + V_o \end{aligned}.$$

The color shifting is accomplished by using an RF carrier frequency, ω_m, for each of the signal voltages, and encoding the communications information into the signal is accomplished with the time varying function $a(t)$. In this configuration the output of the device is a primary lobe of optical energy shifted in frequency from the input energy by ω_m, and side lobes of energy at frequencies shifted by other multiples of ω_m. If configured correctly the efficiency of the double Mach-Zehnder device can be quite high so that nearly half of the input energy is transferred to the output in the primary frequency lobe at $\nu_o + \omega_o$, where ν_o is the frequency of the reference light injected into the input. In this configuration the power in the side lobes is also quite low (approximately –30dB of the power in the primary lobe). This is illustrated in FIG. 11 which shows a plot of the optical power at the

output of the devices versus frequency. Enough spectral range is included to see the primary lobe and its nearest side lobes. The power in the primary lobe at $\nu_o + \omega_o$ referenced to the input power is -3.8 dB. The power in the next highest lobe at $\nu_o - 3\omega_o$ is -30 dB.

Magic Square Software

One of the main difficulties associated with all-optical networks is the need for DWDM colors to be changed as they pass through the core switching network. The need comes from traditional switching architectures such as Clos-type to "scramble" or rearrange subchannels in order to achieve full connectivity.

FIG. 12 shows a schematic of a basic 3-layer Clos switch. Full non-blocking connectivity is achieved between $N*M$ inputs to $N*M$ outputs with the number of switch elements being

$$2*N*M^2 + M*N^2 = M*N*(2*M + N),$$

rather than

$$(M*N)^2$$

for a fully populated switch. This is a large savings if M and N are large. Many-layer Clos switches can achieve even better performance.

The drawback of switching architectures such as the one above is the need to "scramble" the signals in between the crosspoint switches. In the case of the core network, N roughly corresponds to the various fibers coming into a node, and M roughly corresponds to the individual DWDM colors. In order to switch fiber-colors arbitrarily, in the prior art, the system must be able to convert the colors.

In the case of SONET, it is pointless to avoid color conversion, as the fiber-colors are (usually) converted to electrical before they are switched anyway. This process is called OEO (for optical-electrical-optical). Since the switched electrical signals are converted back to light, they might as well have been produced at a different color to start with.

All-optical networks typically require a similar step, but rather than using OEO conversion, most architectures use some kind of non-linear optical wavelength converter, which in many cases is more difficult than a conventional OEO!

In order to *greatly* increase network capability, an all-optical core should avoid color conversion altogether. The trick is how to do this.

Avoiding Color Conversion

If there were a way to assign colors so that provisioning were unnecessary, the core network could be very simple, and thus could handle enormous data rates. The major components would consist of

- Conventional fiber amplifiers
- MEMS-type optical switches
- DWDM splitters/combiners
- No color changing.

Burden is shifted to other parts of the network, most notably the provisioning algorithm, but also onto the system modems:

1. The provisioning algorithm must now take the *connection requirements*, which is the bandwidth requested between every pair of nodes, and convert this into *fiber-color assignments*, which are the colors and fibers that are assigned to connect each pair of nodes. This task is mathematically complex, and it is not immediately obvious that there is even a solution. The difficulty is that, as colors are allocated, certain colors want to be used by multiple links as the other colors are used up. Eventually (without the proper algorithm) the colors cannot be properly allocated, creating a "collision". The algorithm to solve this problem is discussed in the next section.

2. The link assignment between two core nodes or "area codes" takes place over certain defined colors that are generated by the provisioning algorithm. The modems to generate and receive the signals must operate at these assigned DWDM frequencies. This requires a frequency tuning capability on the modems. This frequency-agile modem is described in a following section entitled "Frequency Agile Modem."

Solving these two problems allows bandwidths to the home to be increased by orders of magnitude, as it removes the bottlenecks lying between users.

Solving The Allocation Problem

The provisioning algorithm must take the input *connection requirements*, which is the number of DWDM channels needed between each node, and output *fiber-color assignments*, that is, a color and path for each one of the DWDM requests. We call the *connection requirements* the "requirements matrix" or "R-matrix", and output an "allocation matrix" or "A-matrix".

This algorithm is difficult because, if performed in an unsophisticated manner, certain colors "run out" when they are needed by too many links. This causes "collisions" which are unacceptable solutions.

To get a feel for the problem, we will present a simple hypothetical example. Consider the following R-matrix:

R-matrix	San Diego	Washington	New York	Chicago
San Diego	4	1	2	3
Washington	3	1	2	4
New York	2	1	4	3
Chicago	1	7	2	0

There are a total of 10 colors available to be allocated. The table shows, for example, that Chicago desires 7 colors to transmit data to Washington. The trick now is to assign the 10

colors (0 through 9) so that each city is assigned each color exactly once for both receiving and transmitting.

First, allocate color 0 as follows:

Color 0	San Diego	Washington	New York	Chicago
San Diego		0		
Washington	0			
New York				0
Chicago			0	

Notice that the color appears exactly once in each row and column.

This leaves a reduced R-matrix for the remaining 9 colors:

R-(1-9)	San Diego	Washington	New York	Chicago
San Diego	4	0	2	3
Washington	2	1	2	4
New York	2	1	4	2
Chicago	1	7	1	0

Notice that now each row and column adds to 9. Continuing in this matter, it is a fairly simple matter to allocate the rest of the colors by inspection:

A-matrix	San Diego	Washington	New York	Chicago
San Diego	1,2,3,4	0	7,9	5,6,8
Washington	0,7,8	9	5,6	1,2,3,4
New York	5,6	8	1,2,3,4	0,7,9
Chicago	9	1,2,3,4,5,6,7	0,8	

We can see that this is a valid solution because each row and column has each color exactly once, and the number of colors in each box corresponds to the value in the R-matrix.

Now lets try again, but this time we will start to allocate color 0 in a bad way:

Color 0	San Diego	Washington	New York	Chicago
San Diego			0	
Washington		0		
New York	0			
Chicago				???

Notice that our first three allocations force color 0 into a 4th box that needs no color, and so is a bad allocation. It may seem obvious in this 4x4 example that the first choices were bad; but remember that there are only 4! = 24 permutations of each color. A human can easily "eyeball" all 24 permutations to find a good one. Once the number of nodes gets larger than 10 or so, it becomes nearly impossible for an unaided human to allocate the entire matrix, as each color now has 10! = 3628800 permutations. The problem seems to be most difficult when there are only a few colors remaining.

An algorithm to solve this problem was developed by applicants. We call it the "Magic Matrix" algorithm, because the matrices are similar to the mathematical puzzle "magic squares". The input is an R-matrix such as that shown in FIG. 13.

Here there are 20 input and output channels, with the number of colors requested represented by the intensity, ranging from 0 to 15. Each node is limited to 100 colors total for input (row) and output (column).

The algorithm iteratively assigns wavelengths in such a way as to avoid collisions:

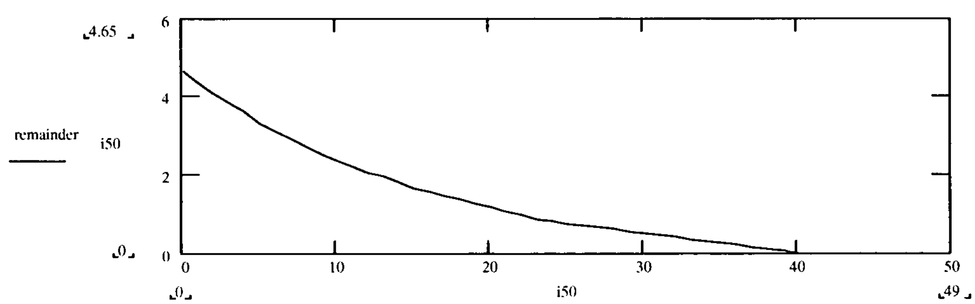

The graph of FIG. 14 shows the residual as a function of iteration number, which becomes zero (finished) at the 39th iteration. This performance is quite good considering that 100 colors are being assigned in those 39 iterations.

Each of the 20x20 squares is now divided into 0-15 colors (depending on the required number of colors from the R-matrix), with each of 100 colors represented as shown in FIG. 15. Each of these 100 colors appears exactly once in each row and column.

The algorithm runs in polynomial time (this means that the number of operations scales as a power of the number of inputs). This is very significant, as the number of arrangements a color can take is the number of permutations of the inputs to the outputs, which is $N!$. This is clearly much too large to sift through, and algorithm intelligence is required to make smart selections without trying every case. The algorithm computational speed is fast enough to service a kilohertz re-provisioning rate.

The algorithm has been tested 50 times. It converged every time in roughly 40 iterations. 5 of those times, however, the algorithm "stalled". When this occurred, the matrix was transposed, and the algorithm converged without a problem.

Our hypothesis on the behavior of the algorithm is that it starts a bad permutation approximately 10% of the time as it approaches the end. This sticking seems to behave randomly, so apparently any permutation or transposition gives the algorithm another 90% chance to succeed. So far the there has never been a case that needed more than a single transpose.

Example

By reference to FIG. 16, the following is an example communication using the network described above.

The caller in San Diego is connected electrically to a local modem 57, sends a request to the modem for a link to Washington as shown at 59. The modem forwards this request to the San Diego area code computer to receive a subchannel on one of the FiberColors assigned to Washington. Here, 2 FiberColors have already been given to the San Diego to Washington connection due to the recent traffic requirements. The San Diego computer picks an available frequency range for the subchannel, and informs the modem at 58, the Washington computer, and the appropriate modem in Washington. The modems then have a direct link and data can be transferred. The Washington modem electrically sends the data to the called party nearby.

Before this happened, due to demands, San Diego had requested from the core allocation computer 2 FiberColors for the link to Washington as shown at 62. The core computer picked the two FiberColors and the link path using the Magic algorithm and other network allocation software. It then instructed the core switches to set up the link. As conditions change, the FiberColors will be reallocated on a sub-second time scale, and the number of FiberColors may go up or down as demands increase or drop.

While the present invention was described in terms of specific embodiment, the reader should understand that many changes and modifications could be made within the general spirit of the invention. For example, the number of area code nodes might be smaller than 250 and the number of frequencies could be less than 300, a small system might have only 20 area code nodes and use only 100 wavelengths, therefore the scope of the invention should be obtained by the appended claims.